Collaborative Search Planning for Multiple Vehicles in Nonhomogeneous Environments

Thomas A. Wettergren and John G. Baylog
Naval Undersea Warfare Center
1176 Howell Street
Newport, RI 02841
t.a.wettergren@ieee.org
john.baylog@navy.mil

***Abstract*—The planning of search paths of multiple unmanned undersea vehicles is complicated by taking into consideration the effects that potential cross-vehicle collaboration may have on their performance. Such collaboration is expected to have the effects of increasing the quality and accuracy of search over that which is obtainable from independent search paths. These problems are important in the undersea domain, where limited communications and uncertain performance characterizations provide a challenge to conduct meaningful search operations in a reasonable amount of time. In this paper, we develop a modeling framework that accounts for the effects of collaborating searchers in a spatially variable undersea environment. The impact of this modeling framework on search plan generation and evaluation is then illustrated using simulation examples.**

I. Introduction

The development of improved methods for performing a search for hidden objects has a long history, going back to the early work of Koopman [1] and others in the 1940s. The early work focused on problems of finding single randomly placed objects in a region using random search patterns. Over time, the scope of problems studied advanced to examining structured ladder search patterns [2] and optimal spatial allocations of search effort [3]. Unfortunately, most of these traditional approaches consider the search for objects that both (a) are positioned according to a uniform distribution over the search region, and (b) are found by a single type of sensor. However, in many undersea applications, searches are conducted for objects that present themselves in a manner that is best suited to being found by one of many available sensors. For instance, the ideal sensor for finding a bottomed object is often very different from the ideal sensor for finding an object suspended in the water column. Because of this, it is desirable to search with sensors tailored to the object of interest. When there exists a mixture of objects, there is expected benefit to having multiple searchers examine the region, each with a sensor tuned to optimally finding a particular one of the objects. Since each sensor still may find the other objects (although not as well as the "tuned" sensor), the use of multi-platform collaboration may benefit the overall search, if it can be properly accounted for at the planning stages.

Recent advances in search evaluation technology allow for collaboration to be assessed at a planning level. These advances treat the search as a probabilistic phenomenon with appropriate joint probabilities used to represent the impact of collaboration. The use of advanced computers with extended memory allows this probabilistic representation to be computed over a dynamic search grid that represents the arbitrary nature of complex search plans. In practice, the search space becomes a representation of search cells, and each search trajectory corresponds to a visit sequence amongst these cells. The collaboration between searchers provides a probabilistic gain for searching a cell that has been previously visited by another searcher. With modern computational memory, it is now feasible to use these cell-based representations of search performance to evaluate competing representative search plans at the planning stages. This allows the consideration of non-traditional search plans and allows for the potential of future optimal search configurations.

The application of search theory in undersea search operations for the purpose of finding a number of sought objects often results in coverage type search paths that are executed by the search platforms. The coverage search trajectories span the search space by an exclusive partitioning with assignment to non-interacting search assets. This is due, in part, to a tendency to assume uniformity in underlying conditions and makes no assumptions on search object placement preference or variation in searcher sensor performance. However, search performance modeling can degrade when a significant spatial variability does exists in either placement preference and/or sensor characteristics, yet uniform conditions are assumed in the planning. The issue of variability is compounded when ancillary search object dependencies are considered that are heterogeneous to spatial characterizations, yet serve to alter the likelihood of detection events. In such cases, not properly accounting for intrinsic variability can lead to incorrect search predictions, which greatly limits the utility of the prediction in the planning process.

By combining a numerical evaluation capability for collaborative search with a framework for vehicle motion planning, we demonstrate a unique capability for accounting for the collaboration in the planning stage of a search. This is shown to provide increased performance over traditional search strategies, and also provides an opportunity to take advantage of any known nonhomogeneity in the search environment. Such nonhomogeneity typically takes the form of multiple search

0-933957-38-1 ©2009 MTS

object classes with searchers specifically tailored to each, or simple geometric preference of search object location due to bathymetric features. With the impact of both collaboration and nonhomogeneity accounted for in the planning process, the full benefit of any collaborative behavior can be achieved. This benefit illustrates the potential impact of various levels of multi-searcher data fusion, allowing search plans that are tailored to each type of fusion to be employed when that fusion is available.

In this paper we develop the methodology for assessing the performance of the search for a mixture of objects by a set of unmanned sensing vehicles that are each tuned to best find a particular object. In the next section, we outline the components of this assessment strategy and provide an overall framework for the search evaluation process. In section III, we analyze the consequences of repeated passes in the evaluation of predicted search performance and the impact that variability has on performance prediction. In particular, we present a method for modeling group search performance which takes into account the additive contributions of the various search sensors as different vehicles pass over the same region (although not simultaneously). Finally, we conclude the paper with a simulation example of a search for a mixture of two types of randomly-placed objects by two searchers that are each tuned to optimally perform against one of the types. This example illustrates the effectiveness of search plans that incorporate multiple search pass collaboration in this search. We summarize the example with a discussion of how these methods may be applied to optimizing the planning for these complex multi-vehicle search operations.

II. Framework for Search Planning, Evaluation, and Adaptation

When planning searches for unmanned vehicles, search theoretic methods use performance estimates of the search sensors to determine search trajectories and any required overlap necessary to achieve the desired search effectiveness. However, existing search evaluation methodologies are based on single searchers and uniform placement distributions. As such, when employed for multiple searchers, there is a natural tendency to perform a prior separation of the search domain into individual sub-domains for each searcher (the so-called "asset allocation" problem). Then collaboration between searchers becomes nothing more than an opportunistic processing strategy, as opposed to planning to maximize the potential collaborative search benefit. Furthermore, when searcher collaboration is ignored in the planning stages, the anticipated performance can be misleading and lead to less efficient search trajectories for the autonomous platform in the given period of time. When trying to utilize scarce resources in a cost-efficient manner, such inefficiencies should necessarily be examined in the context of the overall system design.

The overall search planning, updating, evaluation, and re-planning methodology is outlined in figure 1. In this process, the search problem is described by the two components on the left side of the figure. These two components enter into the *search evaluation* process inside of the upper feedback loop, corresponding to the *plan optimization* process. At each iteration within the plan optimization process, the *vehicle path* is updated and the evaluation is re-processed. Given an optimal plan, we begin execution of the search. As some data is returned from early sorties (represented as the *execute partial search* block), the priors that describe the search object parameterization are updated in the *update priors* process. This lower feedback process provides the ability to re-plan searches based on these new priors, leading to a new iteration of the upper feedback process of plan optimization. This process repeats until the search effort is completed (or at least until the last sortie is planned). In a previous paper [4], we describe the process of updating the priors that is shown in the *update priors* block and the *search evaluation* block. In the next section, we examine the impact of spatial variability on the evaluation of search effort as applied by varying asset types. The *execute partial search* block is not discussed, since it merely represents a physical process (albeit one we simulate in the examples), and the *plan optimization* function is beyond the limited scope of this paper, its development is the subject of ongoing research.

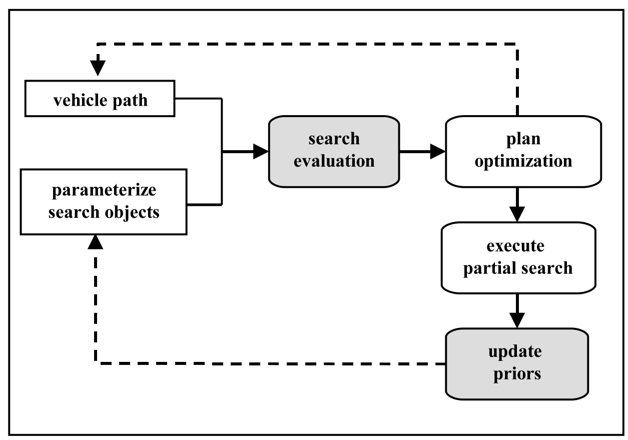

Fig. 1. Unmanned Vehicle Search Planning and Re-planning Process

III. Impact of Variability on Cell-Based Search Performance Evaluation

To mathematically model the search performance of multiple collaborative searchers in nonhomogeneous environments, we employ a decomposition of the search process over a geometric grid. The decomposition considers the separation of search effectiveness into those calculations that relate directly to the target of search via a *contact model*, and those that relate to the searcher via a *search effort model*. When properly decomposed, the combined calculation yields a probability of successful search: that is, the probability of finding what you are looking for. This decomposition process holds as long as the contact model can be factored appropriately; in general, that implies a model of a single object placement preference, whereby it is assumed that the object exists somewhere in the search space. Both the contact model and search effort

model are applied on a common geometric grid whose scale is fine enough to incorporate necessary details of geometric variability (in both contact and search effort), yet coarse enough to include the important search features of the search object in a single cell. This decomposition provides an ability to translate the kinematics of each individual searcher into a sequence of cell pass visits on the search grid. Then all computations can be efficiently applied on a cell basis. We note that the modeling framework implies the search paths are given (and fixed); any adjustment or optimization of the search paths is performed in an operation that wraps around the search evaluation process, as shown in figure 1.

A. Generic Search with Overlapping Scans

Cell based evaluation of search performance under the condition of independent search passes has been well studied in the literature [3], [5]. The absence of any exploitable collaborative dependency allows the recursion in detection likelihood to be based solely upon the single search pass expected probability of detection. An arbitrary level of assurance that the specific object of concern is not present within any given cell can be attained by merely ensuring that enough search passes are executed over that cell. Let $P_{\eta_i}(n)$ denote the probability that an object in cell i has not been detected after n independent search passes. The resulting probability takes the form

$$P_{\eta_i}(n) = (1 - P_{D_i})^n \tag{1}$$

which follows from the joint occurrence of n independent non-detection events. Let $P_{\delta_i}(n)$ denote the corresponding first detection probability afforded to pass n whereby δ follows a *geometric* distribution with probability at the i-th cell

$$P_{\delta_i}(n) = P_{D_i}(1 - P_{D_i})^{n-1}. \tag{2}$$

From these, we let $P_{\Sigma_i}(n)$ denote the aggregate detection probability after n complete independent search passes with

$$P_{\Sigma_i}(n) = \sum_{k=1}^{n} P_{\delta_i}(n) = 1 - (1 - P_{D_i})^n = 1 - P_{\eta_i}(n). \tag{3}$$

Here, we see that the latter form represents the closed form solution to the finite geometric series presented in the summation.

These forms can be rendered sequentially to facilitate the evaluation of alternate search plans. A single recursion in $P_{\eta_i}(n)$ is required for each cell in the search space. Specifically, when search pass detection events are assumed to be independent, we have the joint non-detection event recursion

$$P_{\eta_i}(n) = (1 - P_{D_i})P_{\eta_i}(n-1) \tag{4}$$

with the corresponding incremental pass detection probability

$$P_{\delta_i}(n) = P_{D_i}P_{\eta_i}(n-1) \tag{5}$$

for first detection occurring on the n-th pass.

Let P_C denote a clearance threshold on search effort. When applied identically to each search cell, P_C defines a lower bound on the multi-pass aggregate probability of detection, $P_{\Sigma_i}(n)$ to be achieved in that cell by the search plan. Hence,

$$P_{\Sigma_i}(n) = 1 - (1 - P_{D_i})^n \geq P_C \tag{6}$$

and the number of search passes over that cell must satisfy

$$(1 - P_{D_i})^{n_i} \leq 1 - P_C. \tag{7}$$

Taking the log of both sides of the inequality, rearranging and acknowledging the negative range of log likelihood values yields the lower bound

$$n_i \geq \frac{\ln(1 - P_C)}{\ln(1 - P_{D_i})}. \tag{8}$$

Spatial variability in detection likelihood may cause the number of search passes necessary over the respective cells to vary over the search space as well. Depending upon the degree of variability, imposing this constraint may result in significant expenditure of resource for only incremental gain and may or may not be practical. Alternately, the P_C threshold can be applied collectively to the partitioned search space whereby detection likelihood is marginalized according to placement probability

$$\sum_i P_{\Sigma_i}(n)P_{G_i} = 1 - \sum_i (1 - P_{D_i})^{n_i} P_{G_i} \geq P_C. \tag{9}$$

This form of clearance calculation weights more heavily those cells deemed more likely to have object placement occur within them. Observe that when variability occurs only over placement preference, then $P_{D_i} = P_D$ is constant over the space. For this case, when a coverage strategy is employed such that $n_i = n$ for all i, then equation (9) reduces to

$$\sum_i (1 - P_D)^{n_i} P_{G_i} = (1 - P_D)^n \leq 1 - P_C. \tag{10}$$

and the lower bound of equation (8) applies.

For non-coverage search strategies, spatial variability can induce a variation in the number of search passes conducted over the cells spanning the space. The calculation of the minimum number of cell visits necessary to satisfy equation (9) may not be apparent nor may the solution be unique. However, in order for a deviation from a coverage policy to apply, then the variability in spatial conditions must be significant enough to cause a reordering in the ideal cell visit sequence. Let $\{\pi_k\}_{k=1}^{\infty}$ denote an ideal sequence of cell pass visits where $\pi_k = P_{\delta_i}(n)P_{G_i}$ with $k = k(i, n)$ denoting a one-to-one mapping from $(i \times n) \rightarrow k$ such that $\pi_{k_a} \geq \pi_{k_b}$ for $k_a < k_b$. If k is ordered according to n, then no cell is searched again until all cells have been searched to the same pass level. Otherwise, the ideal cell search pass sequence can become non-monotonic in n with certain cells given multi-pass search preference over other cells. Then, for $n_a > n_b$ given that $k_a(i_a, n_a) < k_b(i_b, n_b)$, we have

$$P_{D_{i_a}}(1 - P_{D_{i_a}})^{n_a - 1} P_{G_{i_a}} \geq P_{D_{i_b}}(1 - P_{D_{i_b}})^{n_b - 1} P_{G_{i_b}} \tag{11}$$

To grasp the significance of equation (11), observe that for the limiting cases when either $P_{D_{i_b}} = 0$ or $P_{G_{i_b}} = 0$, no

search effort is warranted over the cell and any other cell which is non-zero in these quantities will support an infinite number of cell visits in lieu of the unproductive search. For the definitive search, $P_{D_{i_a}} = 1$, only one search pass is necessary and cell ordering will be based solely on geographic likelihood. More generally, let $n_a > n_a - 1 \geq n_b$ be applied to set the upper bound $n_a - 1 = n_b$ on the latter. Let

$$\rho_{ba} = \frac{P_{D_{i_b}} P_{G_{i_b}}}{P_{D_{i_a}} P_{G_{i_a}}} \tag{12}$$

denote a single pass variability ratio for the two given cells. Then, from equation (11), the number of pass sequence reorderings induced by variability in the ideal search pass sequence is finite for each cell pair combination and is bounded by

$$n_b \leq \frac{ln(\rho_{ba}) - ln(1 - P_{D_{i_b}})}{ln(1 - P_{D_{i_a}}) - ln(1 - P_{D_{i_b}})} \tag{13}$$

In particular, observe that the limiting occurrence $n_b = 1$ represents a least restrictive condition and for this case

$$P_{D_{i_a}} \leq 1 - \rho_{ba}. \tag{14}$$

We see that ideal search pass sequences that deviate from coverage type search trajectories occur when the joint likelihood reduction due to spatial variability is large relative to the operating values of detection likelihood. Conversely, the lower the single pass variability ratio, ρ_{ba} becomes, the larger will be the supporting values of detection likelihood that induce search pass sequence reordering.

This notion of cell ordering applies equally well to one or more search vehicles. That is, the cell sequence can be executed by a single vehicle or, more generally, by a coordination of search paths for multiple vehicles. The latter constitutes a *pass-level* collaboration.

B. Search in the Presence of Ancillary Dependencies

We now expand our notion of variability to include functional dependencies that are non-spatial in nature. In particular, we apply the concept of an additive utility function to the recursions of equations (4) and (5) as developed in [6]. These utility functions encapsulate the functional interdependence between the detection likelihood function and ancillary random variables, θ, given the disposition of ancillary scan variables, φ. In [6] we highlighted the continuous random variable describing object orientation relative to the scan variable sensor reference axis. In this paper we emphasize a discrete variation in object type and its corresponding impact on sensor detection performance for one or more sets of searcher characteristics.

For each cell *i*, let $\theta = \{m_j\}_{j=1}^{N_m}$ denote a set of N_m possible object types occurring within the cell with cell mixture probability $P_i(m_j)$. Let $\varphi = \{v_s\}_{s=1}^{N_v}$ denote the non-random set of N_v possible searcher asset types over which to plan the search. Finally, let $P_{D_i}(m_j; v_s)$ denote the cell detection likelihood conditioned on object type and given the searcher type executing the search pass. Here, detection events are considered to be conditionally independent. Hence, individual cell recursions of the form of equations (4) and (5) can be developed. These take the form

$$P_{\delta_i}(m_j; n) = P_{D_i}(m_j; v_{s_n}))P_{\eta_i}(m_j; n-1), \tag{15}$$

and

$$P_{\eta_i}(m_j; n) = (1 - P_{D_i}(m_j; v_{s_n}))P_{\eta_i}(m_j; n-1), \tag{16}$$

where v_{s_n} denotes the searcher conducting the n-th search pass. Marginalizing over object type yields the cell incremental and non-detection probabilities

$$P_{\delta_i}(n) = \sum_{m_j \in \theta} P_{\delta_i}(m_j; n) P_i(m_j) \tag{17}$$

$$P_{\eta_i}(n) = \sum_{m_j \in \theta} P_{\eta_i}(m_j; n) P_i(m_j) \tag{18}$$

In [6], the additive utility functions are developed from first principles. In this paper, we will accept the functional form and work back backwards to explore the impact of variability. The cell incremental search pass probability and non-detection probability recursion that incorporate dependency modeling via utility functions are shown below.

$$P_{\delta_i}(n) = \overline{P}_{D_i}(v_s) P_{\eta_i}(n-1) + U_i(n) \tag{19}$$

$$P_{\eta_i}(n) = (1 - \overline{P}_{D_i}(v_s)) P_{\eta_i}(n-1) - U_i(n). \tag{20}$$

Here, $\overline{P}_{D_i}(v_s)$ represents a reference likelihood level about which variability in detection likelihood is measured. This reference can be developed globally over the search space or individually for each cell as was done in [6]. While either form will produce identical search performance assessments, prudent selection of the reference can reduce the storage and computational time requirements of the utility functions.

For our purposes, let $\overline{P}_{D_i}(v_s) = P_{D_i}(v_s)$ indicate the cell referenced mean detection probability and $\overline{P}_{D_i}(v_s) = P_D(v_s)$ indicate the global reference where

$$P_{D_i}(v_s) = \sum_{m_j \in \theta} P_{D_i}(m_j; v_s) P_i(m_j) \tag{21}$$

$$P_D(v_s) = \sum_i \sum_{m_j \in \theta} P_{D_i}(m_j; v_s) P_{G_i}(m_j) P(m_j). \tag{22}$$

Here $P(m_j)$ denotes static object type mixture weights that apply over the search space and $P_i(m_j)$ denotes the cell conditional mixture weights, which can vary over the space. With the composite placement preference P_{G_i} developed over the multiple object types given by

$$P_{G_i} = \sum_{m_j \in \theta} P(m_j) P_{G_i}(m_j), \tag{23}$$

the latter calculates as

$$P_i(m_j) = \frac{P(m_j) P_{G_i}(m_j)}{P_{G_i}}. \tag{24}$$

To generate the functional form of the utility function, we set equal equations (17) and (19) and combine with (15) to obtain

$$U_i(n) = \sum_{m_j\in\theta} P_{D_i}(m_j; v_{s_n})) \cdot P_{\eta_i}(m_j; n-1) P_i(m_j) - \overline{P}_{D_i}(v_s) P_{\eta_i}(n-1). \quad (25)$$

Naturally, the form of the utility function depends upon the reference level applied. When the cell referenced mean detection probability of equation (21) is used, the utility function takes the form

$$U_i(n) = \sum_{m_j\in\theta} P_{D_i}(m_j; v_{s_n})) P_i(m_j) \times [P_{\eta_i}(m_j; n-1) - P_{\eta_i}(n-1)] \quad (26)$$

Assuming all non-detection probabilities initialize at $P_{\eta_i}(0) = P_{\eta_i}(m_j; 0) = 0$, the cell utility for the first two passes become

$$U_i(1) = 0, \quad (27)$$

and

$$U_i(2) = -\sum_{m_j\in\theta} P_{D_i}(m_j; v_{s_n}))\Delta P_{D_i}(m_j; v_{s_n}) P_i(m_j), \quad (28)$$

respectively, where $\Delta P_{D_i}(m_j; v_{s_n}) = P_{D_i}(m_j; v_{s_n}) - \overline{P}_{D_i}(v_s)$ denotes the perturbation of the object type detection likelihood from the cell mean value. Subsequent terms follow the development in [6] and are functions of the perturbation moments. For this insular reference, all spatial variability in detection likelihood and placement preference must vanish before the cell utility generalizes to a form applicable over the search space. Note, however, that even for this case, search planning may induce variability at the cell level due to variability in the assets with which the space is searched.

When the globally referenced mean detection probability of equation (22) is used, the form of the utility function becomes

$$U_i(n) = \sum_{m_j\in\theta} P(m_j) \times \left[\frac{P_{D_i}(m_j; v_{s_n}) P_{G_i}(m_j)}{P_{G_i}} P_{\eta_i}(m_j; n-1) - \overline{P}_D(m_j; v_s) P_{\eta_i}(n-1) \right] \quad (29)$$

where

$$\overline{P}_D(m_j; v_s) = \sum_i P_{D_i}(m_j; v_s) P_{G_i}(m_j) \quad (30)$$

denotes the grid mean or spatial average of each object type detection probability. Defining

$$\rho_{im}(v_{s_n}) = \frac{P_{D_i}(m_j; v_{s_n}) P_{G_i}(m_j)}{\overline{P}_D(m_j; v_s) P_{G_i}} \quad (31)$$

as a variability indicator relative to the object specific spatial averages, equation (29) reduces to

$$U_i(n) = \sum_{m_j\in\theta} P(m_j) \overline{P}_D(m_j; v_s) \times [\rho_{im}(v_{s_n}) P_{\eta_i}(m_j; n-1) - P_{\eta_i}(n-1)]. \quad (32)$$

Spatial variability will induce a variation in resultant utility function. Consequently, the first pass utility function becomes non-zero and relates the sensor specific search capacity to the resulting search effort. When object detection probabilities become constant over the space, variability is induced by object placement preference with respect to object type and searcher characteristics. When object placement preference becomes uniform, utility becomes dependent upon searcher specialization alone and, as above, cell definition corresponds solely to search trajectory constraints. For this case, the form of the utility function is identical to that developed for the cell reference level.

This kind of functional interdependence between the types of objects sought and the asset capabilities brought to bear in order to find objects represents a much higher degree of collaboration between searchers. That is, exactly how search pass sequences are executed in an *asset-level* collaboration can significantly effect the efficiency of the search.

IV. Numerical Examples

In this section, results are presented to demonstrate the variations in expected performance that are induced by variability within the search space. The results are calculated numerically over a cell based search grid comprised of 250,000 cells. The sequential cell-probability aggregation performed by the algorithm is described in detail in [4] and we present only the results here. We define the problem as having two search assets being deployed to find objects of two possible types. The assets operate with complementary sensor detection capabilities as given below in table I for the two object types. These values apply everywhere over the space. Also provided in the table is the mean detection probability marginalized over object type which is applied for examples where ancillary dependency on object is not assumed.

TABLE I
Searcher-Object Type Detection Capability

Search Asset	Object 1	Object 2	Mean Value
1	0.84	0.36	0.60
2	0.36	0.84	0.60

The objects are assumed to be placed in a square search space according to the spatially variable object density maps depicted in figure 2. These density maps are uniform in the vertical dimension and vary according to a Normal shaping distribution in the horizontal direction. For each object type, the expected number of objects in the space is one. As such, these type specific object density maps act as conditional probability density functions. They are summed and normalized to produce the composite object placement prior depicted in figure 3. Hence, a fixed mixture weight of $P(m_j) = 0.5$ applies for either type. Note that while detection probabilities and mixtures are fixed over the space, the cell conditional mixture weights $P_i(m_j)$ are not, as per equation 24.

The two assets conduct search according to one of the two search paradigms depicted below in figure 4. On the left, the

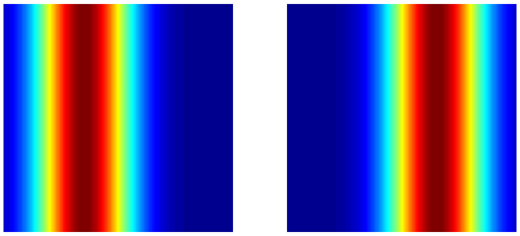

Object Type 1 Object Type 2

Fig. 2. Object Placement Densities

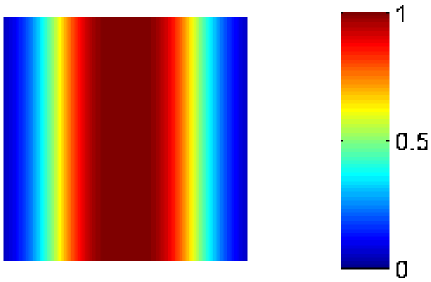

Fig. 3. Composite Object Placement Prior

square search region is divided horizontally into two halves over which vertical ladder searches are conducted. On the right, both assets are deployed over the entire search region and their search paths overlap. To facilitate examination of search pass sequencing effects, the starting point of the search is set to occur at a location where the placement prior drops off to one third of its value at its most likely locations. From this point, a partial search of the space is conducted in the direction of the most likely positions followed by a complete sweep of the space.

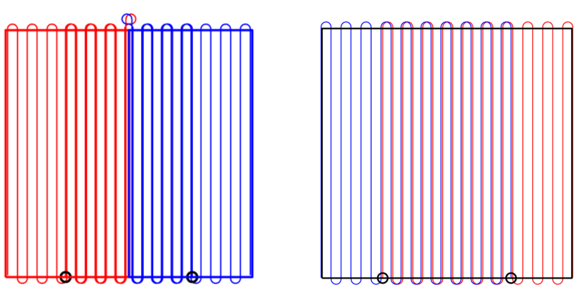

Split Region Search Paths Overlap Search Paths

Fig. 4. Search Path Definitions

We start with the case where no ancillary dependencies are assumed and both search assets operate with the mean detection probability $P_D = 0.6$. In figure 5, clearance probability for the left half search is presented as a function of search time for the split region search plan. This is shown as the upper curve in the figure. Also shown as the lower curve in the figure is the result obtained by employing a reciprocal trajectory for the search. For this search trajectory, the positions of the search asset are reversed as a function of time. Any difference in the expediency of the pass sequence between the two search paths will be indicated by a gap between the curves. As the two paths result in precisely the same set coverage, initial and final performance results are identical. As shown in the figure, this difference can be significant.

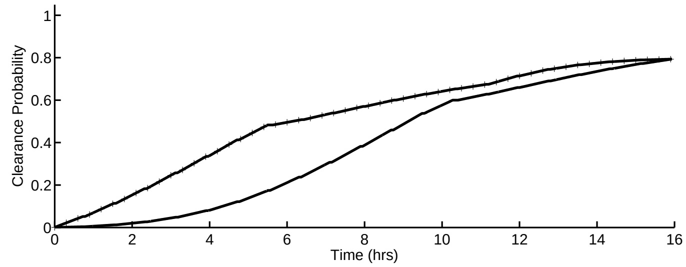

Fig. 5. Impact of Spatial Variability on a Candidate Left Region Search Path and its Reciprocal Trajectory

Next, we compare the effectiveness of the two search strategies for this case of no ancillary dependency. The result is depicted in figure 6. For the split region strategy, clearance probability is calculated individually and summed over the two search regions (shown as the blue curve). For the overlapping search plan, results are calculated and aggregated simultaneously per the multi-platform search evaluation algorithm (shown as the black curve). As variability is restricted to the object neutral placement prior and pass sequencing is reasonably aligned for the two plans, there is hardly any noticeable differences in the resulting performance. The slight difference that does occur is due to the kinematic requirement for the search platform to turn around as part of the split region search policy. This causes a slight delay in execution. For this case, it does not matter which strategy is employed as long as a proper pass sequence is executed.

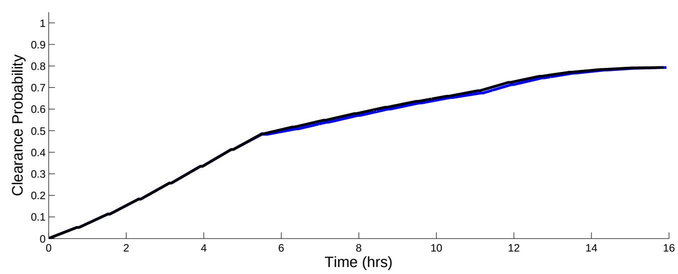

Fig. 6. Search Performance Comparison of Split and Overlap Search Paths

The next two figures concern the inclusion of ancillary dependency as given in the variability exhibited in table I. First we examine the performance for the left region using the split region strategy. Again, results for the search path and its reciprocal trajectory are provided. Clearly, from figure 2, object type 1 is much more likely to occur in this region and one should anticipate good performance when the asset type is matched to the object type. This is shown to occur in figure 7 as the upper red curve corresponding to the selection of asset type 1 as the searcher. The gap between the red curves

highlight the variation in expediency with respect to the search pass sequence. The blue curves denote a selection of asset type 2 as the searcher.

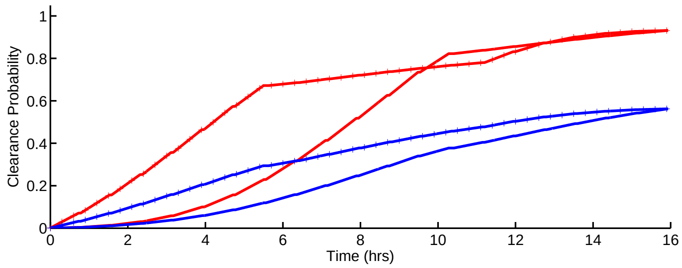

Fig. 7. Impact of Combined Variability on a Candidate Search Path and its Reciprocal Path

A comparison of the overlapping search path trajectories is provided in figure 8. Three performance curves are illustrated. The central black curve corresponds to the case of no ancillary dependence (i.e., $P_D = 0.6$) assumed in the model. It is identical to that shown in figure 6. The red and blue curves depict results for the combined variability case for the same asset selection combination as above. Two interesting observations can be made. First, improved performance over the mean value performance is achieved when ancillary dependence is accounted for in the performance assessment regardless of object-asset search pass coordination. Secondly, exactly which search path exhibits the better performance when this dependency is accounted for depends upon time constraints and clearance level specified. Initially, asset type 1 starting on the left dominates with its matched detection priority until well past the overlap. Then, the more thorough matched-search presented by starting asset 2 on the left compensates for the lack of initial timeliness.

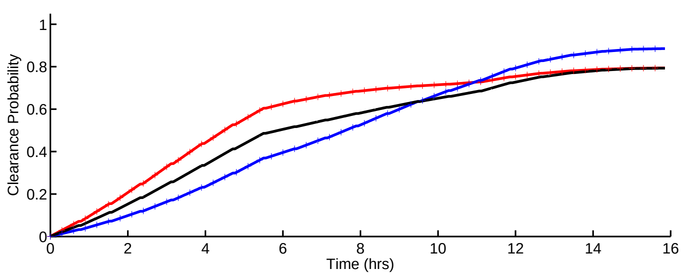

Fig. 8. Search Performance for Overlapping Search Paths

Note that nearly identical performance is achieved for these search trajectories if asset type 1 is selected as the left region searcher using the split region plan and asset type 2 is selected as the left starting searcher in the overlap plan. Both present reasonable search plans provided that the match between region and asset type is exploited. Blind asset allocation can produce the poor results indicated by the blue curves of figure 7. Ultimately, ideal search plans are not unique and are subject to the specific criteria being applied and the knowledge brought to bear in developing the plans.

V. Conclusion

We have presented a methodology and mathematical modeling approach to planning searches for multiple unmanned vehicles. The approach developed is applicable to vehicles searching for objects whose expected location is given by a spatially-varying likelihood function. Furthermore, the approach accounts for collaboration between the searchers, properly accounting for the expected benefit of any planned fusion strategy between searchers. The procedures developed readily apply to planning search paths based on area coverage asset allocations as well as to coordinated group searches.

This approach was applied to a problem of multiple unmanned vehicles searching for objects whose location is unknown, but are expected to appear according to a known spatial likelihood pattern. Furthermore, mixtures of various different object types are allowed, each with a potentially unique likelihood pattern. In a simulation setting, we illustrated how the benefits of a group of different searchers, each tailored to a specific object from the mixture, can be improved via their collaborative behavior, as opposed to performing a simple non-collaborative sub-division of the search region. In this manner, we have illustrated a capability to interpret spatial variability in placement, sensor/object specificity, and multi-searcher collaboration.

The example presented in this paper represents a single realization of a general tool set for evaluating search of multiple collaborating vehicles in the context of placement uncertainty. Future extensions of this work include the addition of expected errors (such as navigational uncertainty) and the use of these methods in a search plan optimization setting.

Acknowledgment

This work was sponsored by the Office of Naval Research Code 321MS.

References

[1] B. H. Koopman, *Search and Screening: General Principles with Historical Applications*, Pergamon Press, Elmsford, NY, 1980.
[2] J. C. Hyland, C. M. Smith, and P. J. Stampnick, "Multi-dimensional MCM Tactical Analysis and Optimal Resource Allocation", CSS/TR-95/13, Naval Surface Warfare Center, Panama City, FL, February, 1993.
[3] L. D. Stone, *Theory of Optimal Search*, 2nd Edition, ORSA Books, Arlington, VA, 1989.
[4] J. G. Baylog, T. A. Wettergren, J. C. Hyland, and C. M. Smith, "Robust search for structured object placement using unmanned vehicles," in *Proc. OCEANS 2008*, Oct. 2008.
[5] S. Benkoski, M. Monticino, and J. Weisinger, "A survey of the search theory literature," *Naval Research Logistics*, vol. 38, pp. 469–494, 1991.
[6] T. A. Wettergren and J. G. Baylog, "Modeling sequential searches with ancillary target dependencies," manuscript under journal review.